DOCUMENT RESUME

ED 206 652 TM 810 542

AUTHOR Ludlow, Larry H.
TITLE An Exploratory Investigation of Rasch Model Residuals.
PUB DATE Apr 81
NOTE 37p.: Paper presented at the Annual Meeting of the American Educational Research Association (65th, Los Angeles, CA, April 13-17, 1981).

EDRS PRICE MF01/PC02 Plus Postage.
DESCRIPTORS *Goodness of Fit; *Latent Trait Theory; *Models; Statistical Analysis
IDENTIFIERS *Data Interpretation; Rasch Model; *Residuals (Statistics)

ABSTRACT

Residual patterns should be studied in order to understand when and why data deviate from a model. This paper illustrates some techniques for exploring residual patterns resulting from the difference between observed and expected scores as predicted by a Rasch model for polychotomous data. The score residual divided by its standard deviation is called a standardized residual. A variety of plots of residuals are presented to illustrate residual pattern interpretation. A systematic analysis of residuals can offer the investigator a decision facilitation technique not found in conventional summary fit statistics. (BW)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED206652

U.S. DEPARTMENT OF EDUCATION
NATIONAL INSTITUTE OF EDUCATION
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

☒ This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official NIE position or policy.

AN EXPLORATORY INVESTIGATION OF RASCH MODEL RESIDUALS

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY

L. H. Ludlow

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

Larry H. Ludlow

MESA
Department of Education
University of Chicago

and
Rehabilitative Engineering Research
and Development Center
Hines Veterans Administration
Hines, Illinois

Paper presented at the American Educational Research Association
Annual Convention, Los Angeles, Cal., April 12-17, 1981

TM 810 542

Models are convenient expressions of how we think things ought to be. When we use models to understand experience, however, it is inevitable that less than a perfect explanation results. This need not mean the end of the model. By studying the residuals from our expectations we can learn from, and come to, a better understanding of our experiences.

The importance of inspecting residuals to gain information leading to better models and better understood data has been clearly demonstrated in the analysis of variance literature by, among others, Anscombe and Tukey (1963) and Draper and Smith (1966). Likewise, the analysis of covariance structures depends on the interpretion of residuals in the model testing stages(Joreskog,1979). The psychometric literature, however, does not reveal as comprehensive an effort at the investigation of residual patterns in order to understand when and why data deviate from a model. For example, when the Rasch model is applied to dichotomous data the best analyses may compute fit statistics for items across people and for people across items (Mead, 1975). But these statistics are often inadequate to locate the source of model departure. And, since these statistics are sensitive to sample size and test length there is controversy regarding the magnitude to be regarded as misfitting.

Although these summary statistics have demonstrated their usefulness, the analysis of model fit must be carried further. Just as Draper and Smith urge the display of residual patterns in

addition to summary statistics , so may the same approach be applied in our work. This paper illustrates some techniques found useful for exploring residual patterns resulting from the difference between observed and expected scores as predicted by a Rasch model for polychotomous data developed by Geoff Masters (Masters, 1980; Masters and Wright, 1981).

We can derive three alternative forms with which to compute a residual from the model. The observed minus expected score yields what we call a score residual. The score residual divided by its variance is called a logit residual. The score residual divided by its standard deviation is called a standardized residual.

Our choice of the residual to inspect was based on how well it provided useful information. The score residual was rejected because floor and ceiling effects restricted the variation in the extremes. This led to our formulation of the logit residual which proved to be sensitive to variations in the extremes, but perhaps too much so. The same patterns could be found when comparing logit and standardized results but the adjustment to tables and pictures in order to handle the magnitude of logit residuals proved to be an inconvenience. Since the standardized residuals manifest a familiar metric, interpretation of their patterns became the easiest of the choices.

In Sloane's (1981) discussion the summary approach to residuals suggested some misfitting items. Positive fits result when

more able children score lower than expected (yielding negative residuals) and less able children score higher than expected (yielding positive residuals). Negative fits result when able children do better than expected (positive residuals) and less able children do worse than expected (negative residuals). But these summary statistics do not tell us exactly who has scored unusually on which items, nor how widespread the problem is.

The first picture created to study this was a plot of standardized residuals against child abilities (Figure 1). This plot contains the original sample of 500 children. From left to right the abilities increase. X's represent 10 or more children with a given ability and residual. Each column of points shows the spread of residuals for the children with that ability. For instance, the leftmost column shows the 14 residuals for the least able child. The picture we expect for data fitting the model is a random pattern with a mean near zero and a standard deviation near one for each vertical array of residuals.

What we see in Figure 1 indicates something else. A rectangle marking off plus and minus 3 standard deviations was added to the picture to highlight the asymmetry of the distribution in the 4th quadrant. The points outside the area are able children with large negative residuals. It is recognized that very able children will appear unusual only when they miss easy items, hence large negative residuals. Likewise, the least able children will only appear unusual when they succeed on relatively hard items,

yielding large positive residuals. The task, then, becomes one of finding out who the able children are, which easy items were missed, and whether a valid explanation, other than one of chance occurence, can be suggested for the large negative residuals.

One way to see what is happening is to print the person by item residual matrix and to scan it for patterns. This matrix can be constructed according to various sorting schemes depending upon what aspect the investigator wishes to highlight. It may also be built according to group membership. Such a matrix may indicate unusual behaviors individuals or groups of people have had on single items or blocks of items. Examples of this type of matrix can be found in Wright and Stone (1979).

Figure 2 shows such a matrix. In this example the columns are items sorted by their difficulty and the rows are children sorted by their ability. Both sorts are descending. Various summaries may be collected in the margins. The entries in the matrix are standardized truncated residuals. Any residual with an absolute value equal to or greater than 9 was set equal to 9. The summary statistics, however, were computed from real valued standardized residuals. In these matrices, a criterion must be set so that data are sufficiently focused when person by item interactions are brought out. This allows for the matrix to be simple enough to grasp. This criterion was set such that only negative residuals equal to or less than -3.0 and only children who had at least one such residual are displayed. This matrix concentrates on

surprising mistakes.

In the SUBJ. ID field in the left margin of the matrix, the first 2 columns are the age of the child in months and the third column is the child's sex. ("1" for boys, "2" for girls). The right most margin has the child's logit ability.

Inspection of the ID field suggests that boys and girls are spread equally throughout the ability range. There does appear to be some age differentiation, however. The older children generally have the greater abilities while the younger children are primarily in the lower ability range. The fact that older children perform better on a developmental instrument is hardly surprising. Of interest to us is whether the failures on relatively easy items can be explained by a characteristic of the item that led to a related group of children having had unexpected trouble with it.

In order to simplify this demonstration we chose to emphasize what appeared to be an age factor operating in an unexpected direction. First we controlled for sex by looking only at boys and then selected the 56 oldest and 42 youngest ones.

Figure 3 is the plot of the 98 boy abilities by their residuals. This plot corresponds with that of Figure 1. At this point, though, we do not know the age of the most able children in Figure 3, nor what items are involved, nor if there is even a dif-

ference in the abilities distributed among the two age groups.

Figure 4 presents the residual matrix for the reduced set of 98 boys. In the ID field the first 2 columns are the age and the third column is the age group classification ("1" if the age is less than 47 months, "2" if the age is greater than 59 months). The criterion was set for negative residuals equal to or less than -2.0. This matrix highlights those more able children who have done worse than expected.

This concludes our discussion of search techniques for understanding the summary statistics used by Sloane. These results do not completely explain her misfitting items because the girls are not included as part of our older and younger boy dichotomy. The point should be clear, however, that the misfit in Sloane's original solution can be explored through a residual analysis. Inspection of the residual matrices is one way of understanding why an item has misfit the model.

Now we look at the sub-sample of 98 boys to see if the variable definition has remained the same for the two age groups. Our intent is to demonstrate a process that may be applied whenever the question is asked "Have these groups of people performed the same on the instrument?".

Figure 5 is an ability frequency distribution map. The youngest boys are on the upper map, the older boys on the lower. The

numbers indicate the number of people located at the same position. Double digit numbers are read vertically, e.g. in the second map the 1 above the 8 refers to 18 people with an ability near 3.0. The M and S refer to the respective means and standard deviations. Figure 5 shows how much more able the older boys are than the younger. Figures 6 and 7 show the ability by residual plots for the younger and older boys, respectively. There were no negative abilities for the older boys so we removed the negative side of the plot from Figure 7. A comparison with Figure 3 confirms that the large negative residuals are due primarily to the older boys. We now turn to the investigation of the items.

Figure 8 is a plot of residuals against item difficulties. The items range from easiest at the left to hardest at the right. The rectangle marks off plus and minus 3 standard deviations. The 3rd quadrant shows that the easier items have the large negative residuals. This picture confims what we learned from Figures 3 and 4.

The information in the sorted matrices of Figure's 2 and 4 is convenient when the sample in the matrix does not exceed 100. At that point three pages are required to present the picture. This motivates us to seek another way of presenting the residual information for large samples or when group comparisons are to be made. A useful way to compare groups in terms of residual distributions over individual items is shown in Figures 9 and 10. In Figure 9, for ITEM8, we see the residual distribution for each

group of boys on a line extending from -5 to +5 standardized residual units. The integers on the line represent the number of boys who had a given residual value. Again the numbers are read vertically and the M and S represent the group means and standard deviations of these residuals.

This type of map allows a detailed examination of within and between group performances on individual items. Also, we can combine items that are similar and build these maps for groups of people across items. Not only can we find where surprising behavior has occurred but we can identify which children and items are involved.

A useful summary of all the item maps is contained in Table 1. The Table contains the group means and standard deviations for each item's residuals. This table may be used in a variety of ways. We could take the difference between group means and look at the items with the largest differences. We could look at each item where there was a difference in the sign of the mean. For this example, we chose to concentrate on the 2nd column of means and standard deviations, in particular, the negative means.

When data fit the model, residuals are ability and difficulty free and differences in summary statistics should be attributable to chance. We have seen in these data, however, a combination of high ability and low difficulty that results in high negative residuals. Since the older boys are generally the more able we

expect them to have higher means and smaller standard deviations in the residuals relative to the younger ones. The negative means for the older boys under column 2 represent instances where the mean performance by the older boys was worse in terms of expected behavior than that expected from the younger ones. Since this instrument was designed to measure development it is surprising to see some of the more able, older boys performing considerably worse than expected. This is not to say their performance was bad- just that it was not as good as was expected.

We need to know if the surprising residual means can be explained by chance or whether something more fundamental is at stake. From inspecting the maps for ITEM8 and ITEM14, in Figure 9, it is evident that outliers (indicated by asterisks) have skewed the means of the older boys. These children may also be found in Figure 4 under the two items in question. After tracing the problem down to two outliers on each item, we conclude the items are functioning as intended but we could take the matter further if we cared to because we have identified specific children who had surprising trouble.

Looking at the residual distributions for ITEM6, ITEM7 and ITEM11, in Figure 10, we see that in each case the older boy distribution is bimodal with a group of children about one standard deviation below the mean. Those with positive residuals are not surprising. Those with the negative residuals suggest a pattern of deficiency or erroneous observation and we may ask if the

larger, negative residuals are from the same children. From a matrix of residuals with a criterion set at -1 we found that, in fact, some of the same children had large residuals on combinations of two of the three items and these children were in the middle range of the ability distribution for the older boys. Apparently the problem with ITEM6 and ITEM7 is due to a few of the same older boys. Since the median for these items is considerably to the right of the respective means we suggest these items are functioning as intended. It would be reasonable, however, to monitor these items in future applications.

Such an interpretation is not possible, however, for ITEM11 in Figure 10. Here the median approaches the mean but the distribution is still bimodal. Who made the mistakes on this item? Interestingly, it was the most able of the older boys. We determine this by looking again at Figure 4 and noticing the largest residuals belong to boys in the highest ability level. The problem with ITEM11 may be due to confusion on the part of some of the older boys who did not know to whom the 'in front of' instruction was directed.

Where does this leave us? We could be asked why we did not calibrate the items separately for the two groups of boys and plot the item difficulties against one another with error bands. This was done in Figure 11 to underscore why it is not necessary and can be misleading. One reason for exploring the residuals from a total sample calibration was to show that the same group

differences are seen as when separate calibrations are done. The first thing to check in Figure 11 is whether the same items that the residual analysis tagged as peculiar also stand out in the plot. We see they do in fact either lie outside or close to the 3 standard error bands. The important point, however, is not the corroboration. If we saw only this plot we might conclude there was a variable definition problem in these data, one where the items meant one thing to the older boys and another to the younger. But this is not the case as we have seen by exploring and explaining the residual patterns. The plot in Figure 11 merely indicates that certain items are peculiar but does not explain why. The same argument holds true if one is interested in computing t-statistcs to test the difference between pairs of item difficulties. The residual analysis not only identifies the same items but may explain with judicious investigation why there was a peculiarity and shows that, except for ITEM11, it is not necessarily a problem of item construction.

In conclusion, we arque that a systematic analysis of residuals offers the investigator a decision facilitation technique not found in the conventional summary fit statistics. The process may be used to understand individual people, individual items, groups of people or groups of items. Such an understanding of residual patterns may prove useful as a means of addressing issues of 'item bias', 'quessing', and 'discrimination'.

References

Anscombe,F.J. and Tukey,S.W. "The examination and analysis of residuals," Technometrics,5,141-160 (1963).

Draper,N.R. and Smith,H. Applied Regression Analysis. John Wiley and Sons,New York, 1966.

Joreskog,K.G. and Sorbom,D. Advances in Factor Analysis and Structural Equation Methods. Abt Books, Cambridge, Mass. 1979.

Masters,G.N. A Rasch model for rating scales. Unpublished doctoral dissertation, University of Chicago, 1980.

Masters,G.N. and Wright,B.D. Rating Scale Analysis. Chicago:MESA Press, 1981.

Mead,R.J. Analysis of fit to the Rasch model. Unpublished doctoral dissertation, University of Chicago, 1975.

Sloane,K.D. "An application of the partial credit model," Paper presented at the American Educational Research Association Annual Convention, Los Angeles, Ca., April, 1981.

Wright,B.D. and Stone,M.H. Best Test Design. Chicago:MESA Press, 1979.

AN EXPLORATORY INVESTIGATION OF RASCH MODEL RESIDUALS

LARRY H. LUDLOW

MESA
DEPARTMENT OF EDUCATION
UNIVERSITY OF CHICAGO

AND

REHABILITATIVE ENGINEERING RESEARCH
AND DEVELOPMENT CENTER
HINES VETERANS ADMINISTRATION
HINES, ILLINOIS

PAPER PRESENTED AT THE AMERICAN EDUCATIONAL RESEARCH ASSOCIATION
ANNUAL CONVENTION, LOS ANGELES, CAL., APRIL 12-17, 1981

(TABLES AND FIGURES)

DIAL: FINE MOTOR, CONCEPT: ALL TODDLERS INCLUDED: STD RESIDUALS

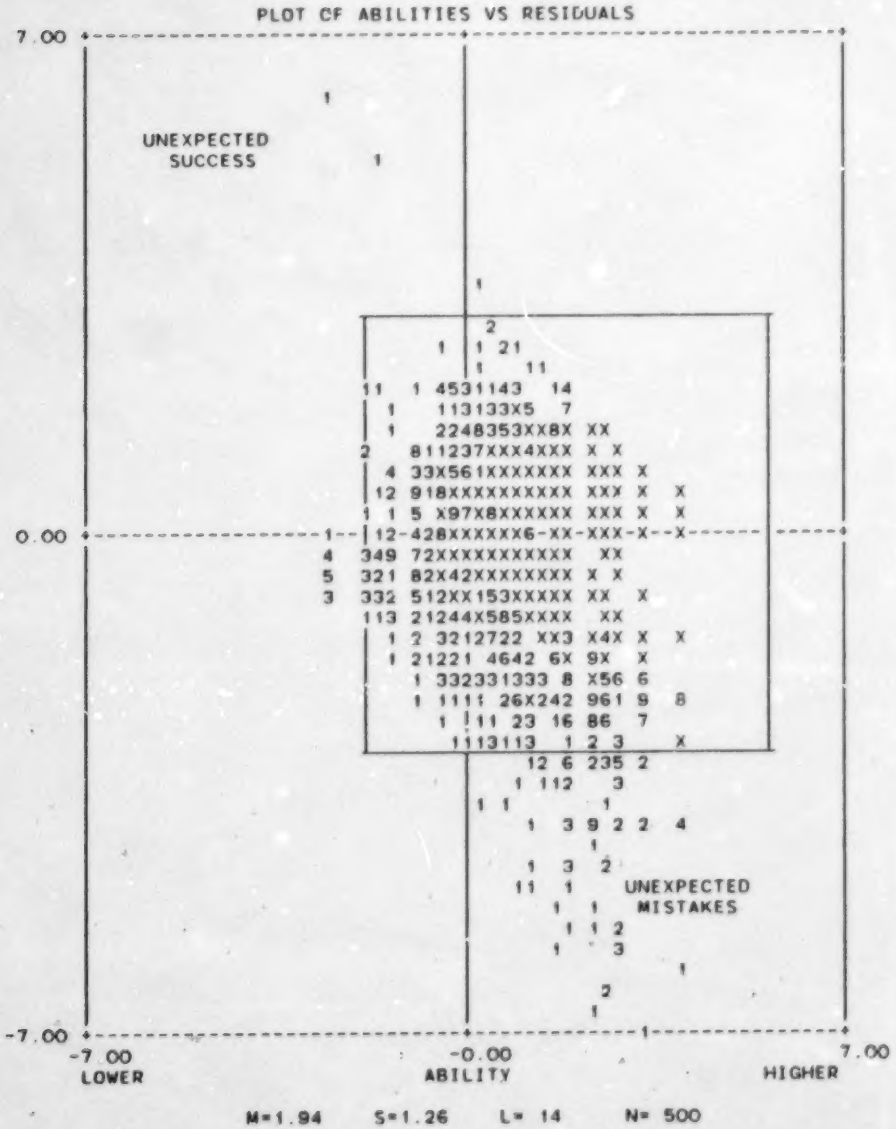

THE RECTANGLE IS CENTERED ON THE MEAN ABILITY FOR THE 500 CHILDREN. THE MEAN (M), STANDARD DEVIATION (S), TEST LENGTH (L) AND SAMPLE SIZE (N) ARE GIVEN BELOW THE PLOT.

M=1.94 S=1.26 L= 14 N= 500

L.H. LUDLOW, MESA PSYCHOMETRIC LABORATORY, UNIVERSITY OF CHICAGO

FIGURE 1

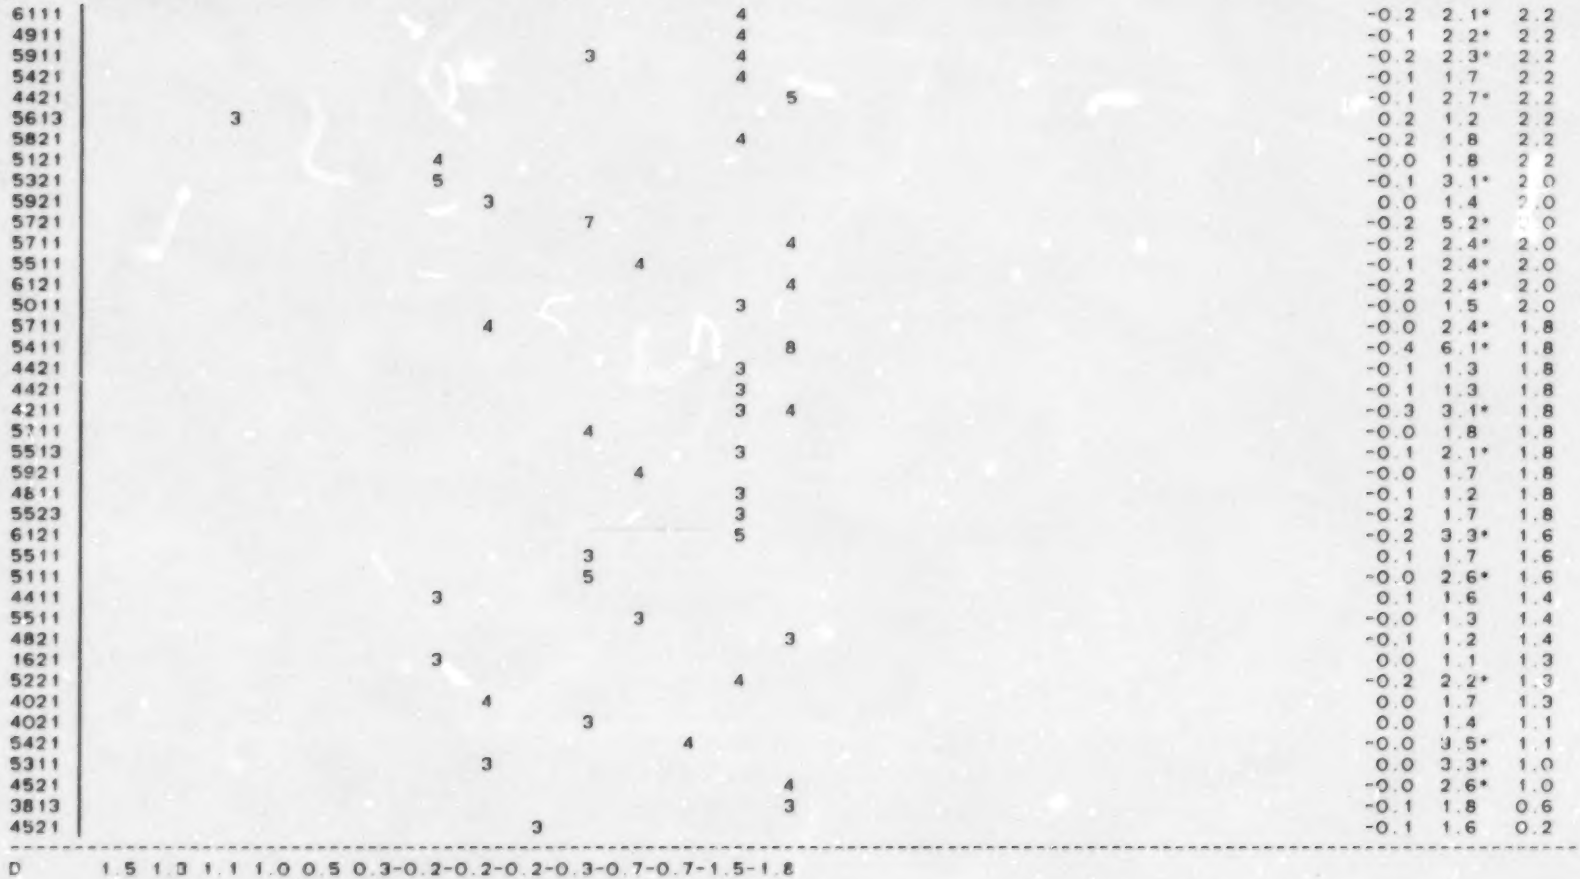

L.H. LUDLOW, MESA PSYCHOMETRIC LABORATORY, UNIVERSITY OF CHICAGO

FIGURE 2

DIAL: FINE MOTOR, CONCEPT: AGE=BOYS LT 47 OR GT 59 MONTHS:STD RESIDUALS

M=1.8 S=.81 L= 14 N= 98

L.H. LUDLOW, MESA PSYCHOMETRIC LABORATORY, UNIVERSITY OF CHICAGO

FIGURE 3

MATRIX OF SORTED TRUNCATED RESIDUALS FOR THE SUB-SAMPLE OF 98 BOYS

COLUMNS ARE ITEMS SORTED BY DIFFICULTY
ROWS ARE PEOPLE SORTED BY ABILITY

RESIDUALS=(OBSERVED-EXPECTED) SCORES
*** ONLY NEGATIVE RESIDUALS LE-2 ARE SHOWN ***

SUBJ. ID	ITEM NAMES #5	#4	#6	#7	#11	#10	#3	#2	#13	#8	#9	#14	#12	#1	MEAN	MS	B
6421												6			-0.2	2.9*	4.1
6021					2										0.1	0.5	4.1
6422									3						-0.0	1.2	4.1
6021									3						-0.0	1.2	4.1
6321						2									0.0	0.7	4.1
6021						2									0.0	0.7	4.1
6221					2										0.1	0.5	4.1
6121					2										0.1	0.5	4.1
6421										7					-0.2	3.7*	3.3
6021			2												0.0	0.8	3.3
6421									2						0.0	0.9	3.3
6221												4			-0.1	1.5	3.3
6021												3			-0.0	1.1	2.9
6421									2			3			-0.1	1.4	2.9
6021												3			-0.0	1.1	2.9
6221												3			-0.0	1.2	2.9
6121				2											0.1	0.8	2.5
6121										3					-0.1	1.7	2.5
3711							2								0.0	0.7	2.2
6121											2		4		-0.2	2.1*	2.2
4211													3	4	-0.3	3.1*	1.8
6121				2											0.1	0.8	1.6
4111	2														-0.1	1.4	1.1
6221											2				-0.0	0.9	1.1
4111														2	-0.1	1.0	1.0
4213											2				0.0	1.1	0.9
3813														3	-0.1	1.8	0.6
4211												2			0.0	1.0	0.6
4111													2		-0.2	1.4	0.5
2511									2						-0.0	0.6	0.4
3711												2			-0.1	1.3	-0.1
3811		2													-0.1	1.4	-0.2
3511									2						-0.2	1.1	-0.9
0	1.5	1.3	1.1	1.0	0.5	0.3	-0.2	-0.2	-0.2	-0.3	-0.7	-0.7	-1.5	-1.8			

NEGATIVE RESIDUALS FROM THE 98 BOYS

L.H. LUDLOW, MESA PSYCHOMETRIC LABORATORY, UNIVERSITY OF CHICAGO

FIGURE 4

FREQUENCY DISTRIBUTION OF 98 PEOPLE ABILITY MEASURES IN 2 GROUPS

BOYS YOUNGER THAN 47 MONTHS
M= 0.30 SD= 0.86 N= 42.

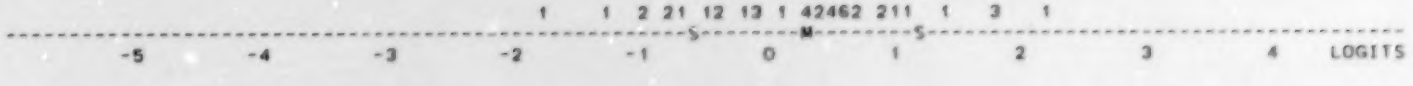

BOYS OLDER THAN 59 MONTHS
M= 2.93 SD= 0.76 N= 56.

L.H. LUDLOW, MESA PSYCHOMETRIC LABORATORY, UNIVERSITY OF CHICAGO

FIGURE 5

PLOTS OF ABILITIES VS STANDARDIZED RESIDUALS

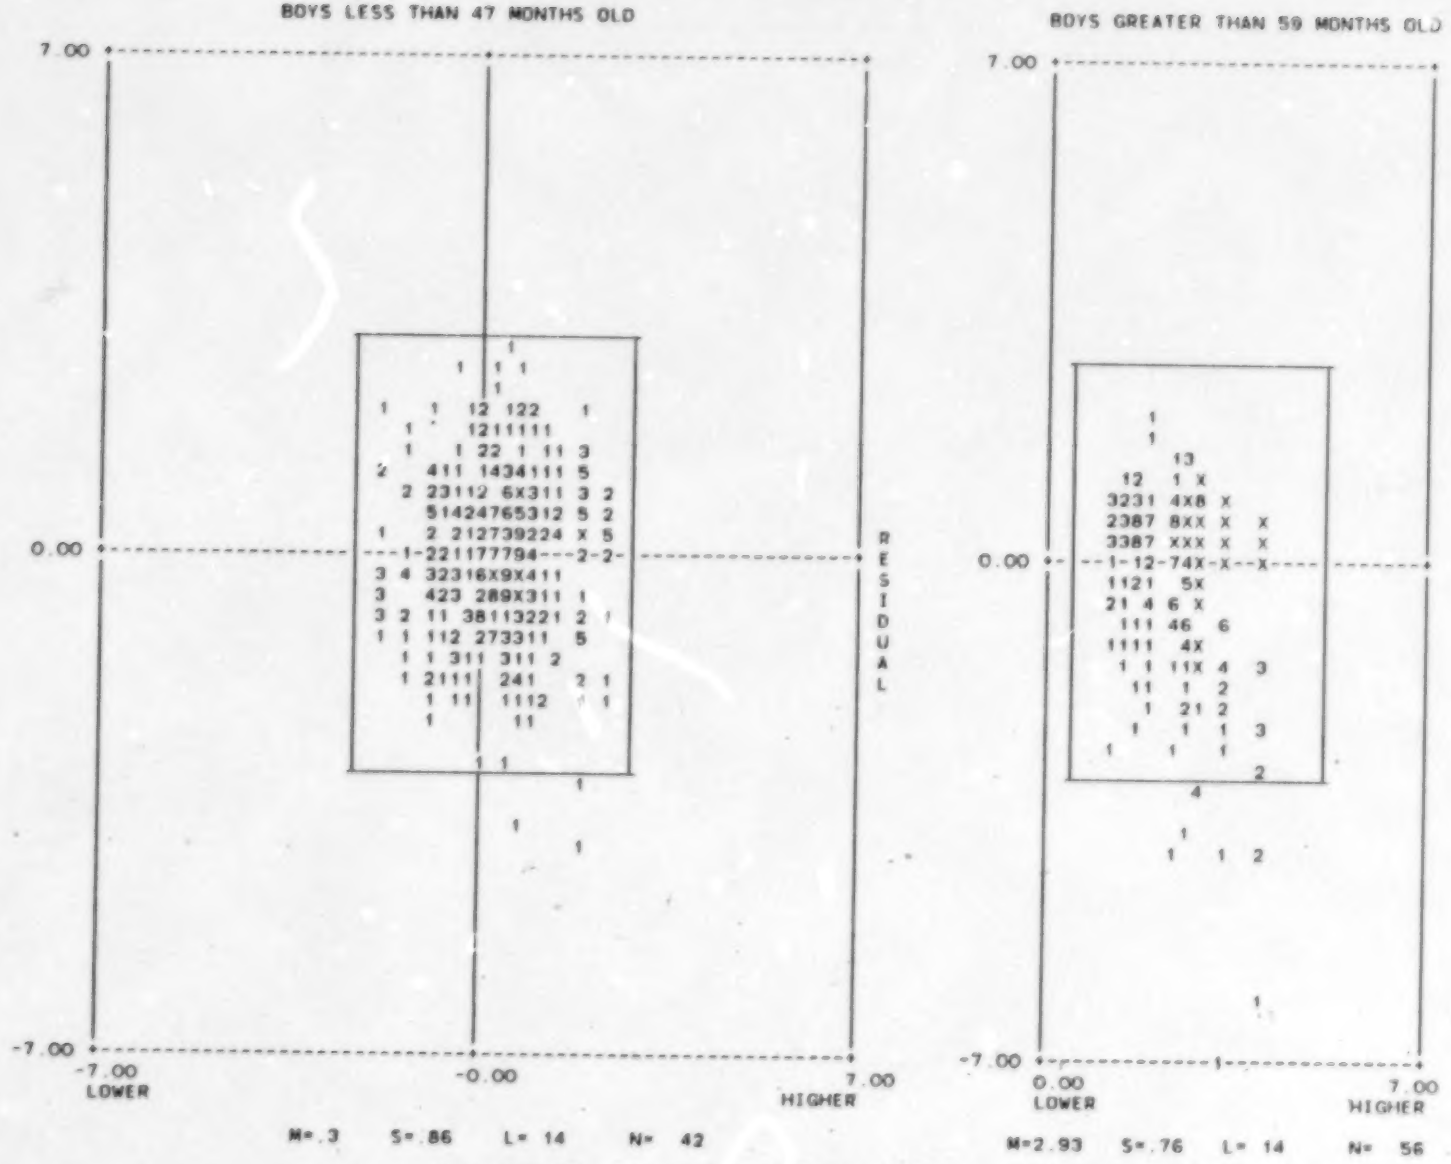

L.H. LUDLOW, MESA PSYCHOMETRIC LABORATORY, UNIVERSITY OF CHICAGO

FIGURE 6

FIGURE 7

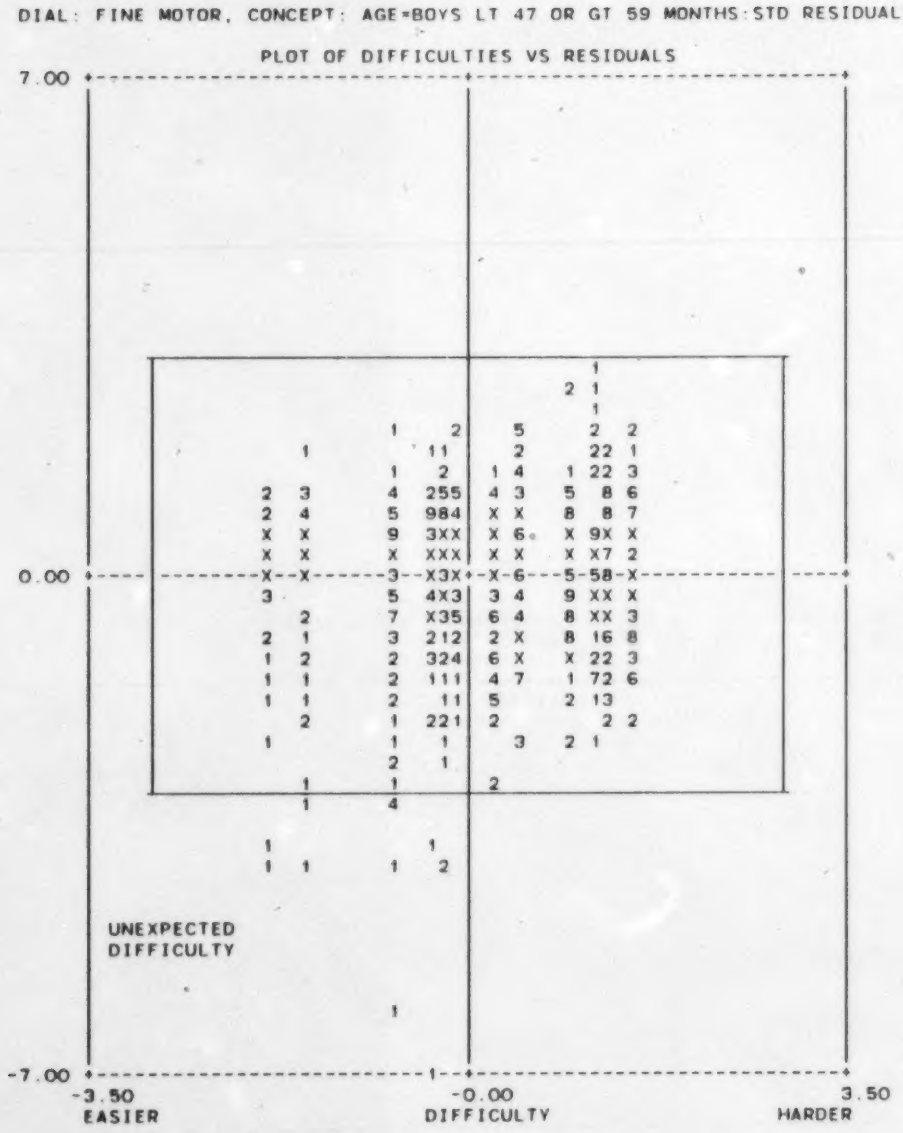
DIAL: FINE MOTOR, CONCEPT: AGE=BOYS LT 47 OR GT 59 MONTHS:STD RESIDUALS
PLOT OF DIFFICULTIES VS RESIDUALS
7.00
0.00
-7.00
RESIDUAL
UNEXPECTED
DIFFICULTY
-3.50
EASIER
-0.00
DIFFICULTY
3.50
HARDER
L= 14 N= 98
L.H. LUDLOW, MESA PSYCHOMETRIC LABORATORY, UNIVERSITY OF CHICAGO

FIGURE 8

STANDARDIZED RESIDUAL FREQUENCY DISTRIBUTIONS

ITEM8

BOYS YOUNGER THAN 47 MONTHS
M= 0.01 SD= 0.80 N= 42.

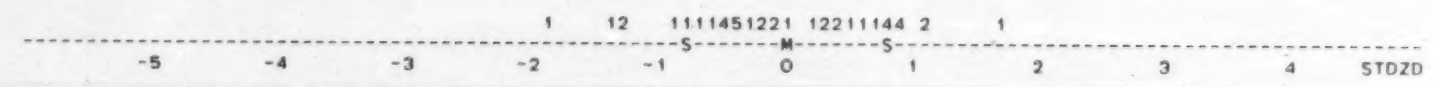

BOYS OLDER THAN 59 MONTHS
M= -0.09 SD= 1.14 N= 56.

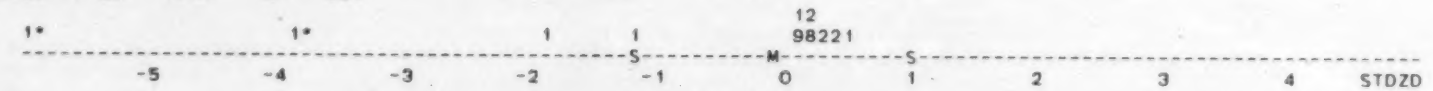

ITEM14

BOYS YOUNGER THAN 47 MONTHS
M= -0.18 SD= 1.00 N= 42.

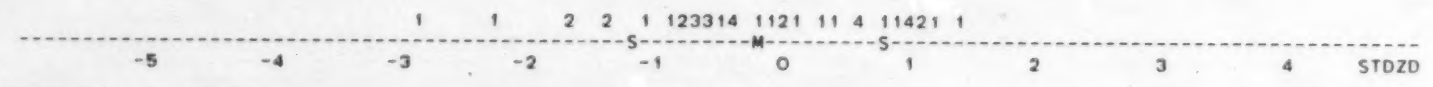

BOYS OLDER THAN 59 MONTHS
M= -0.20 SD= 1.39 N= 56.

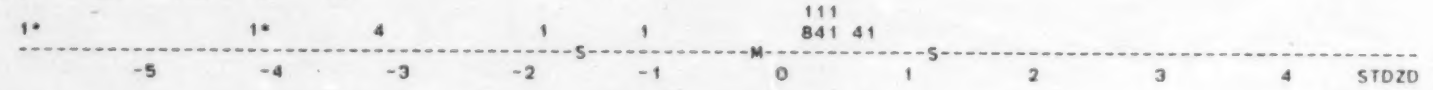

L.H. LUDLOW, MESA PSYCHOMETRIC LABORATORY, UNIVERSITY OF CHICAGO

FIGURE 9

ITEM6

BOYS YOUNGER THAN 47 MONTHS
M= 0.13 SD= 1.06 N= 42.

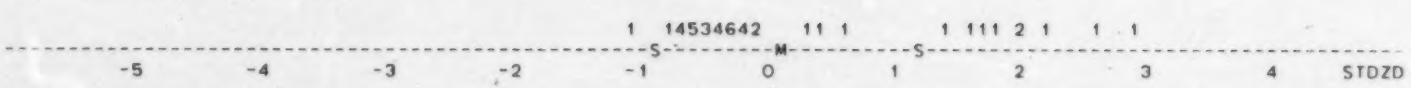

BOYS OLDER THAN 59 MONTHS
M= -0.07 SD= 0.80 N= 56.

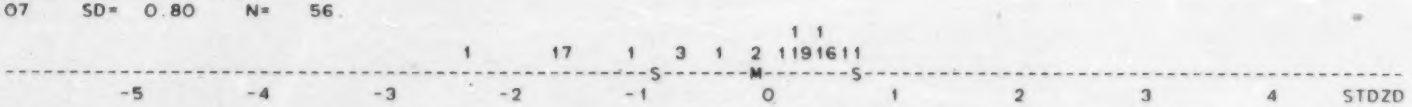

ITEM7

BOYS YOUNGER THAN 47 MONTHS
M= 0.11 SD= 0.94 N= 42.

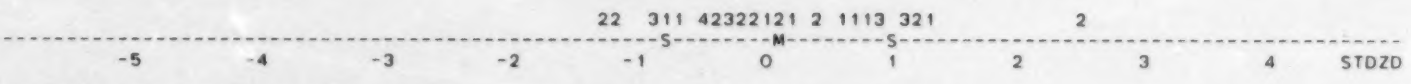

BOYS OLDER THAN 59 MONTHS
M= -0.07 SD= 0.89 N= 56.

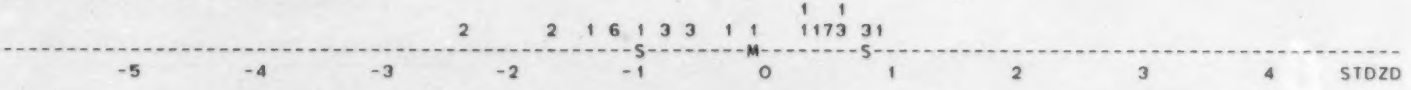

ITEM11

BOYS YOUNGER THAN 47 MONTHS
M= 0.41 SD= 1.03 N= 42.

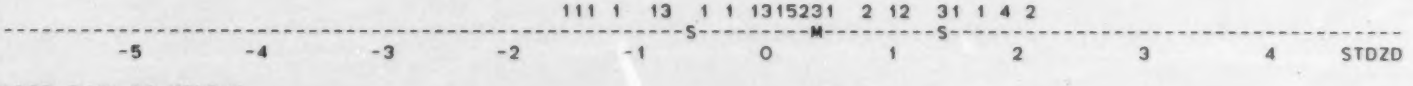

BOYS OLDER THAN 59 MONTHS
M= -0.28 SD= 1.00 N= 56.

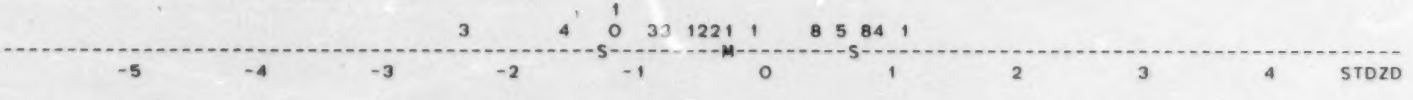

L.H. LUDLOW, MESA PSYCHOMETRIC LABORATORY, UNIVERSITY OF CHICAGO

FIGURE 10

DIAL: FINE MOTOR, CONCEPT: AGE=BOYS LT 47 OR GT 59 MONTHS:STD RESIDUALS

RESIDUAL SUMMARY TABLE:
ITEMS LISTED IN SEQUENCE

GROUPS ACCORDING TO ID CODE

ITEM NAME	D	YOUNGER BOYS 1 MN	YOUNGER BOYS 1 SD	OLDER BOYS 2 MN	OLDER BOYS 2 SD
#1	-1.81	-0.02	1.14	0.13	0.06
#2	-0.19	-0.22	0.85	0.30	0.50
#3	-0.15	0.05	0.96	0.26	0.13
#4	1.27	-0.49	0.70	0.29	0.79
#5	1.47	-0.34	0.84	0.23	0.77
#6	1.13	0.13	1.06	-0.07	0.80
#7	0.97	0.11	0.94	-0.07	0.89
#8	-0.33	0.01	0.80	-0.09	1.14
#9	-0.67	0.00	1.04	0.06	0.62
#10	0.26	-0.04	0.85	0.03	1.01
#11	0.48	0.41	1.03	-0.28	1.00
#12	-1.51	0.05	1.12	0.10	0.57
#13	-0.20	0.06	0.92	0.15	1.03
#14	-0.71	-0.18	1.00	-0.20	1.39

THIS TABLE CAN FOCUS ON, AMONG OTHERS:
1) ABSOLUTE VALUE MEAN DIFFERENCES,
2) ITEMS WITH REVERSED SIGNS IN THE MEANS,
3) ITEMS WHERE THE OLDER BOYS MEAN
IS LESS THAN THE YOUNGER BOYS MEAN.

L.H. LUDLOW, MESA PSYCHOMETRIC LABORATORY, UNIVERSITY OF CHICAGO

TABLE 1

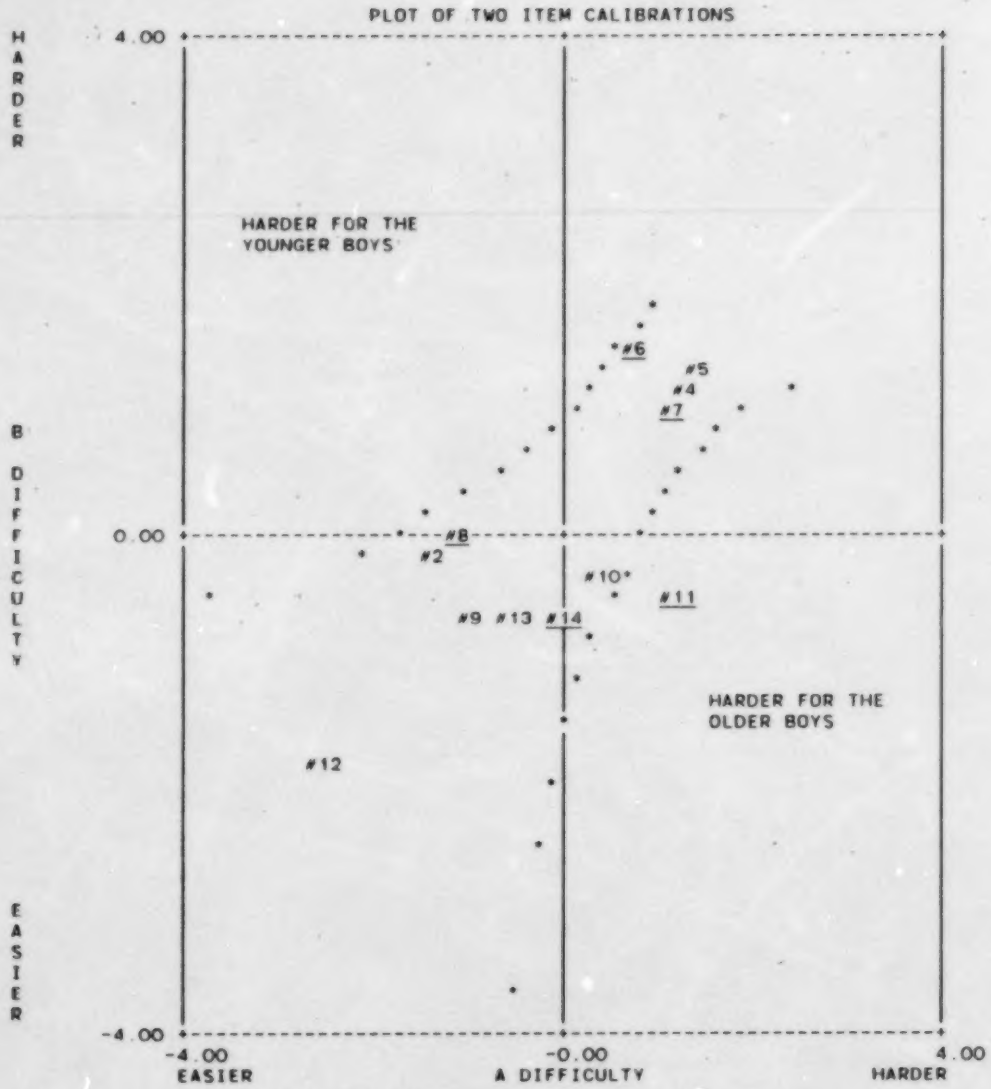

MEAN A= 0.00 SA= 1.28 R=0.677
MEAN B= 0.13 SB= 0.99 N= 12

GROUP A: ITEMS FOR BOYS GREATER THAN 59 MONTHS OLD
GROUP B: ITEMS FOR BOYS LESS THAN 47 MONTHS OLD

THE CONFIDENCE INTERVAL REPRESENTS 3. STANDARD ERRORS
L.H. LUDLOW, MESA PSYCHOMETRIC LABORATORY, UNIVERSITY OF CHICAGO
FIGURE 11

THE UNDERSCORED ITEMS ARE THE SAME UNEXPECTED ITEMS SEEN IN TABLE 1.